

SYSTEM AND METHOD FOR INTERLEAVING INFINIBAND SENDS AND RDMA READ RESPONSES IN A SINGLE RECEIVE QUEUE

5

ABSTRACT

A system and method for processing interleaved Sends of encapsulated communications and responses to RDMA Reads in a single InfiniBand queue pair receive queue. The queue is implemented as one or more linked lists of memory buckets, and stores Send commands (containing encapsulated communications or
10 RDMA Read descriptors for retrieving a communication) until their associated communications are assembled and forwarded to a transmit module. The queue grows as new InfiniBand packets are received, and shrinks as communications (e.g., Ethernet packets) are forwarded. A next packet pointer identifies the next Send command whose communication should be assembled. If it is an
15 encapsulated communication, the communication is forwarded. Otherwise, RDMA Read requests are issued and the responses bypass the tail of the queue and are assembled in an assembly area at the head of the queue.